

CLAIMS

What is claimed is:

1. A method of operating a destination storage server to mirror a primary volume maintained by a source storage server, the method comprising:

receiving a plurality of log entries from the source storage server, the plurality of log entries representing write requests received by the source storage server;

writing the received log entries to a file at a mirror site;

receiving data from the source storage server at a consistency point;

using the data to update a mirror volume at the mirror site via a storage access layer in the destination storage server; and

using log entries from the file to update the mirror volume.

2. A method as recited in claim 1, wherein the storage access layer on the destination storage server implements a RAID protocol, and the mirror volume is a RAID volume.

3. A method as recited in claim 2, wherein said receiving data from the source storage server comprises receiving the data over a network in a protocol which is at a higher logical level than RAID, and wherein the method further comprises passing the data within the destination storage server to the storage access layer after receiving the data from the source storage server.

4. A method as recited in claim 3, wherein said receiving the data from the source storage server comprises receiving the data via TCP/IP.

5. A method as recited in claim 1, wherein said using log entries from the file to update the mirror volume is done at the consistency point in conjunction with said using the data to update the mirror volume.

6. A method as recited in claim 1, wherein said using log entries from the file to update the mirror volume is done in response to a failure that renders the primary volume inaccessible.

7. A method of mirroring data, the method comprising:

at a source storage server, receiving a plurality of write requests from a set of clients;

creating log entries for the write requests in the source storage server;

transmitting the log entries to a destination storage server at a mirror site;

writing the log entries to a file corresponding to the source storage server in the destination storage server; and

at a consistency point,

causing a primary volume maintained by the source storage server to be updated based on the write requests, and

causing a mirror volume maintained by the destination storage server at the mirror site to be updated to reflect the updated primary volume, by

transmitting consistency point data from the source storage server to the destination storage server,

receiving the consistency point data at the destination storage server,

updating the mirror volume through a storage access software layer on the destination storage server based on the received consistency point data, and using log entries from the file to update the mirror volume.

8. A method as recited in claim 7, wherein the primary and mirror volumes are RAID volumes, and wherein the storage access software layer on the destination storage server is a RAID software layer.

9. A method as recited in claim 8, wherein said receiving the consistency point data at the destination storage server comprises receiving the consistency point data over a network in a protocol which is at a higher logical level than RAID, and wherein the method further comprises passing the consistency point data within the destination storage server to the storage access layer after receiving the consistency point data from the source storage server.

10. A method as recited in claim 7, wherein said receiving the consistency point data at the destination storage server comprises receiving the consistency point data via TCP/IP.

11. A method as recited in claim 7, wherein said using log entries from the file to update the mirror volume is done at the consistency point in conjunction with said using the consistency point data to update the mirror volume.

12. A method as recited in claim 7, wherein said using log entries from the file to update the mirror volume is done in response to a failure that renders the primary volume inaccessible.

13. A method of mirroring data, the method comprising:

performing a log forwarding process that includes

- at a source file server, receiving a plurality of write requests from a set of clients,
- creating a log entry for each of the write requests in a nonvolatile memory in the source file server,
- transmitting the log entry for each of the write requests to a destination file server at a mirror site, and
- writing the log entry to a file of log entries corresponding to the source file server in the destination file server;

at a consistency point, performing a data synchronization process that includes

- causing a RAID primary volume maintained by the source file server to be updated based on the write requests, by communicating through a RAID software layer on the source file server, and
- causing a RAID mirror volume maintained by the destination file server at the mirror site to be updated to reflect the updated RAID primary volume, by
- transmitting consistency point data from the source file server to the destination file server using a protocol at a higher logical level than RAID,
- receiving the consistency point data at the destination file server,

and

- updating the RAID mirror volume through a RAID software layer on the destination file server based on the received consistency point data; and
- using log entries from the file to update the RAID mirror volume.

14. A method as recited in claim 13, wherein said transmitting consistency point data from the source file server to the destination file server using a protocol at a higher logical level than RAID comprises transmitting the consistency point data from the source file server to the destination file server using TCP/IP.

15. A method as recited in claim 13, wherein said using log entries from the file to update the RAID mirror volume is done at the consistency point in conjunction with updating the RAID mirror volume based on the received consistency point data.

16. A method as recited in claim 13, wherein said using log entries from the file to update the mirror volume is done in response to a failure that renders the primary volume inaccessible.

17. A storage server comprising:

- a network interface to communicate with a remote storage server;

- a storage access layer to maintain a mirror volume of data reflecting a primary volume of data maintained by the remote storage server;

- a file system to receive a plurality of log entries from the remote storage server, the log entries representing write requests received by the remote storage server from a set of clients, wherein the file system layer further is to write the log entries to a file and, in the event of a failure, to use the file update a mirror volume of the primary volume;
- and

- a network administration unit to receive consistency point data from the remote storage server at the consistency point, and to respond to the received consistency

point data by causing the storage access layer to use the consistency point data to update the mirror volume.

18. A storage server as recited in claim 17, wherein the storage access layer implements a RAID protocol, and the primary and mirror volumes are RAID volumes.

19. A storage server as recited in claim 17, wherein the network administration unit is at the same logical level as the file system layer.

20. A method as recited in claim 17, wherein the consistency point data is received from the remote storage server over a network in a protocol which is at a higher logical level than RAID.

21. A method as recited in claim 20, wherein the consistency point data is received from the remote storage server via TCP/IP.

22. A method as recited in claim 17, wherein the file system uses the file to update the mirror volume at the consistency point in conjunction with using the consistency point data to update the mirror volume.

23. A method as recited in claim 17, wherein the file system uses the file to update the mirror volume in response to a failure that renders the primary volume inaccessible.

24. An apparatus for mirroring data, the apparatus comprising:

means for receiving a plurality of log entries from a remote storage server, the plurality of log entries representing write requests received by the remote storage server;

means for writing the log entries to a file;

means for receiving consistency point data from the remote storage server at a consistency point,

means for using the consistency point data to update a mirror volume through a local storage access software layer, and

means for using the log entries in the file to update the mirror volume.

25. An apparatus as recited in claim 24, wherein the storage access software layer implements a RAID protocol, and the mirror volume is a RAID volume.